

TECHNIQUE FOR CONFIGURING PROCESSORS IN SYSTEM WITH LOGICAL PARTITIONS

Background Of The Invention

5 Field of The Invention

The invention relates to the field of multiprocessing, and in particular, to the distribution of resources on a system where multiple operating system images (same or different) are involved.

10 Background Information

In early computer systems, a single set of program instructions was used to control the basic computer system operations such as input/output, disk access, etc. This set of program instructions was referred to as the 'operating system.' Early computers also generally had

15 only a single central processing unit or CPU where the program instructions were executed. Modern computer systems often have multiple processing units and/or are capable of running multiple operating systems at the same time. Each of the multiple operating systems running on such a modern computer may be confined to operate in a

20 respective logical partition, i.e., an processing space reserved for that operating system. In a shared system, one user may be running the Windows operating system, while another user may be running Unix, for example. Also, there could be multiple users running the same operating system, e.g., Unix, but each running their own copy or 'image' in different 25 respective partitions of the system. These partitions serve to prevent the operating systems running concurrently from interfering with each other.

On a modern computing system running multiple operating system images (i.e., the same operating system or different operating systems) as logical systems in respective partitions, one of the key items to be 30 distributed between each logical system is processing resources. These resources can be dedicated to a particular partition or shared among

multiple partitions. Configuration of these resources takes many user interface forms with GUI (graphical user interface) being the most common.

Even the earliest computer system included a rudimentary mechanism that allowed it to receive input from its users. These mechanism are often called user interfaces. The user interfaces of early computers often employed lights and switches that allowed the computer system user to communicate with the computer system in the system's native language. While this form of communication was effective, its use was essentially limited to scientists and engineers who were trained to understand the particulars of the computer system. For this reason, computer system providers moved to user interfaces that provided for entry of textual commands. While these command-based user interfaces were easier to use than their predecessors, they still typically involved a large number of commands that were often difficult for non-technical individuals to understand and remember. The shortcomings of command-based user interfaces led to the introduction of graphical user interfaces, which are often referred to as GUIs. Instead of requiring the user to have in-depth knowledge of the computer system, graphical user interfaces allow the every day user to communicate with the computer system via a pointing device such as a mouse. To use the computer system to perform various tasks, the user simply "points and clicks" on various GUI items such as icons and taskbars. These icons and taskbars are user friendly ways of activating different programs on the computer.

Techniques for managing the allocation of the processing resources within the user interface generally requires allocating a quantity of the processors. For dedicated processors, the quantity is generally a whole number representing the number of processors. For shared processors, the quantity can be either a numerical value representing both whole and partial processing units (i.e., 2.33 to describe two and a third processors) or a weighting scheme indicating a percentage of the shared

processors (i.e., a partition weight of 100 with all partition weights totaling 1000 for 5 shared processors to indicated a half of a processor). These techniques are user interface independent values which could be applied to a graphical interface, a command line interface, or a menu driven text

5 interface.

The problem with these techniques is that they do not reflect the real life situation of how users think of the processors in their system. System administrators and operators generally think of the processing capacity of their system in a workload measurement value not the number

10 of processors.

There are many known measurements of the processing capacity of the system. One simple measure is MIPS or million instructions per second. Other more sophisticated measures are derived from industry standard workloads such as TPC-C (a specific workload defined by the

15 Transaction Processing Performance Council) and CPW (commercial processing workload). These workloads values are generally well known for server and mainframe class systems and are commonly used to describe system performance in marketing and sales literature.

Configuring logical partitions using CPW values is a technique employed in the V5R1 logical partitioning GUI (graphical user interface). This GUI (a component of the IBM™ AS/400™ Client Access Express for Windows licensed program product), however, will not make suggestions of dedicated versus shared processors.

Therefore, a way to overcome the above problems which would enable system administrators to configure processors for logical partitions using a capacity-based measurement, and which would suggest whether to use dedicated or shared processors, would make the computer systems more user friendly and the system administrators more productive.

30

Summary Of The Invention

It is, therefore, a principal object of this invention to provide a method and apparatus for configuring processors in a system with logical partitions

5 It is another object of the invention to provide a method and apparatus that solves the above mentioned problems so that configuration of logical partition processors is simplified.

These and other objects of the present invention are accomplished by the method and apparatus disclosed herein.

10 According to an aspect of the invention, a capacity-based measurement technique is used to configure processors for logical partitions.

According to an aspect of the invention, instead of specifying a quantity of processing units for a logical partition, the user interface 15 accepts a processing capacity value in units of work, such as CPW or TPC-C, or any well-defined workload measurement units.

According to an aspect of the invention, an exemplary routine for implementing the invention includes allowing a user to select the units to use for configuring the processors, e.g., CPW or TPC-C, or any well-20 defined workload measurement units.

According to an aspect of the invention, the exemplary routine also prompts the user for the quantity of processing power required, i.e., the partition workload capacity required. These two steps can be combined into one step depending on the design of the user interface. That is, the 25 user can be prompted to fill-in the numerical quantity and units on the same screen, or to fill-in the numerical value next to the appropriate units to be used, for example.

According to an aspect of the invention, using a look-up table, for example, the system work capacity for the target system is obtained in the 30 appropriate units from the preceding steps.

According to an aspect of the invention, the number of partition processors is then calculated using the formula the number of partition processors equals the total number of system processors, times the partition workload capacity divided by the system work capacity.

5 According to an aspect of the invention, the resolution for the number of partition processors result is preferably two digits to the right of the decimal, i.e., two decimal places.

10 According to an aspect of the invention, the calculated number of partition processors is then tested to see if it is within a predetermined percentage, e.g., twenty-five percent, of the next full processor increment. If within the predetermined percentage, then the routine recommends using dedicated processors. Otherwise, using shared processors is recommended.

15 According to an aspect of the invention, the calculated number of partition processors and the recommended mode, i.e., shared versus dedicated processors, are displayed to the user so that the user can validate or change the values.

20 According to an aspect of the invention, after validation, the processors are configured according to the settings determined by the routine.

According to an aspect of the invention, a system administrator enters the partition's processing capacity, and the user interface and logical partition management code performs any necessary calculations to determine the number of processing units.

25 According to an aspect of the invention, dedicated or shared processing resources may be recommended based on predetermined criteria.

According to an aspect of the invention, the technique is independent of the type of user interface.

30 According to an aspect of the invention, the configuration of logical partition processors is simplified by utilizing output from the manual

system planning exercise (a workload capacity value) into configuration for the partition directly without any intermediate calculations.

These and other aspects of the invention will become apparent from the detailed description set forth below.

5

Brief Description Of The Drawings

Fig. 1 illustrates a flow chart of a process according to an exemplary embodiment of the present invention.

10 Fig. 2 is a block diagram of a computer system for implementing a technique for configuring processors according to an exemplary embodiment of the invention.

Detailed Description Of The Preferred Embodiment(s)

15 The invention will now be described in more detail by way of example with reference to the embodiment(s) shown in the accompanying figures. It should be kept in mind that the following described embodiment(s) is/are only presented by way of example and should not be construed as limiting the inventive concept to any particular physical configuration.

20 This invention relates to a technique of using a capacity-based measurement to configure processors for logical partitions. Instead of specifying a quantity of processing units for a logical partition, the user interface accepts a processing capacity value in units of work like CPW or TPC-C, or any well-defined workload measurement units.

25 For example, a server system may have 2 processors and be capable of 1000 CPW. The system administrator wants to create a logical partition to support a 250 CPW workload.

To implement this partition using existing user interface techniques, the system administrator would need to utilize the following formula:

30

number of partition processors = (total number of system processors) * (partition workload capacity) / (system work capacity)

This would indicate that 0.5 processing units or half a processor
5 are required.

To implement this partition using the invention technique, the system administrator would enter the partitions processing capacity in CPW as 250. The user interface and logical partition management code would then perform any necessary calculations to determine the number
10 of processing units and could even suggest dedicated or shared processing resources.

This technique is independent of the type of user interface.

An exemplary routine for implementing the invention could take the following form, as illustrated in Fig. 1. The routine initially allows the user
15 to select the units to use for configuring the processors, e.g., CPW or TPC-C, or any well-defined workload measurement units. The routine also prompts the user for the quantity of processing power required, i.e., the partition workload capacity required. These two steps could be combined into one step depending on the design of the user interface.
20 That is, the user could be prompted to fill-in the numerical quantity and units on the same screen, or to fill-in the numerical value next to the appropriate units to be used, for example.

The user selections are limited by what types of units are used in the look-up table. If a new workload standard were to be used, the user
25 interface could be adapted and the table changed.

Using a look-up table, for example, the system work capacity for the target system is obtained in the appropriate units from the preceding steps. The context indicates the model of the target system.

If the target system model is not stored in the database, the user
30 can enter the system work capacity manually, e.g., from sale literature.

For the next calculation, the partition workload capacity and the system workload capacity are expressed using the same units, e.g., both in TPC-C or CPW, or any well-defined workload measurement units.

5 The system work capacity of the target system is stored on the target system or the configuring system.

Systems generally come with this value pre-stored somewhere on-board, e.g., in VPD (vital Product Data) storage, or VPD can enable 'client' lookup, i.e., system model 'xyz' = 123 CPW, for example, model vs. capacity. However, one skilled in the art would understand that there are 10 many different ways in which the system capacity could be stored and accessed.

There is no conversion formula for converting from one unit of processing capacity to another. The table would store capacity values in all the different units needed.

15 The number of partition processors is then calculated using the formula previously described, i.e., the number of partition processors equals the total number of system processors, times the partition workload capacity divided by the system work capacity. The resolution for the result would preferably be two digits to the right of the decimal.

20 Then the calculated number of partition processors is tested to see if it is within a predetermined percentage, e.g., twenty-five percent, of the next full processor increment. If within the predetermined percentage, then the routine will recommend using dedicated processors. For example if the result were 0.80 partition processors then dedicated 25 processors would be recommended. Otherwise, using shared processors would be recommended. The factors that determine the percentage to use as a test for shared processors are based on a definition of what is considered 'reasonable.' A good user interface allows selecting this value under an "Advanced" parameter.

30 The calculated number of partition processors and the recommended mode, i.e., shared versus dedicated processors, is

displayed to the user so that the user can validate or change the values. After validation, the processors are configured according to the settings determined by the routine. This invention simplifies the configuration of logical partition processors.

5 Fig. 2 is a block diagram of a basic computer system 200 for
implementing a technique for configuring processors according to an
exemplary embodiment of the invention. The basic computer system 200
includes partition processing means 201, storage means 202 in which a
program implementing a graphical user interface for configuring
10 processors 203 as described above is disposed, an input device 204, e.g.,
mouse and/or keyboard, and a display means 205.

The invention may be embodied as a computer program product which includes a recording medium, such as a floppy disk, a high-capacity read only memory in the form of an optically read compact disc or CD-

15 ROM, a tape, a transmission type media, such as a digital or analog communications link, or a similar computer program product.

It will be apparent to one skilled in the art the manner of making and using the claimed invention has been adequately disclosed in the above-written description of preferred embodiments taken together with the drawings. The above-described preferred embodiments of the invention are susceptible to modifications, changes, and adaptations, and the same are intended to be comprehended within the meaning and range of equivalents of the appended claims. Although equivalent components may have been mentioned herein which could be used in place of the components illustrated and described with reference to the preferred embodiment(s), this is not meant to be an exhaustive treatment of all the possible equivalents, nor to limit the invention defined by the claims to any particular equivalent or combination thereof. A person skilled in the art would realize that there may be other equivalent components presently known, or to be developed, which could be used within the spirit and scope of the invention defined by the claims.